THE AI BLUEPRINT, EXPLAINED

Decide where AI belongs

Before an association or nonprofit buys AI, or lets it spread further on its own, it should know what staff already use, where AI would earn its place, and what has to be true first. The AI Blueprint puts all of that on paper in four to six weeks, with a policy and a first use your leadership can approve.

01 Why decide first

02 Is it right for you?

03 How the weeks run

04 What you receive

05 What it looks like

06 How organizations use it

For
Anyone planning, building or running your systems

Time to read
About 10 minutes

Best timing
When AI questions start arriving

01 · WHY DECIDE FIRST

Most organizations are already using AI. Very few have decided where it belongs.

AI usually arrives before anyone decides on it: a chat assistant on a personal account, a note-taker someone switched on, a feature that came with a software update.

That is not a failure of discipline. It is what happens when useful tools are easy to reach and nobody has said what is allowed. The decisions that matter come before any product: which uses would earn their place, what happens when each one is wrong, whether the data it would read is ready, and who reviews the result.

We built the AI Blueprint to answer those questions on your side of the table. It uses the same method as our Digital Readiness Blueprint, pointed at a different question, and everything it produces belongs to you. Start the first use with us, with another partner, or on your own. Either way, you start on purpose.

Sunil Dua
Kloudserv

In one paragraph

A fixed-fee engagement of four to six weeks for associations and nonprofits. We find the AI already in use, score 15 to 30 possible uses by value, risk and data readiness, check the data behind the best of them, and hand you a one-page AI policy, a 12-month roadmap and the first use specified, with a decision briefing in your voice.

Quoted before you commit

The fee is fixed. You know it before anything starts.

Vendor-neutral

We recommend the kind of tool each use needs, not a product we resell.

02 · IS IT RIGHT FOR YOU?

It is a good fit if

- Your systems are in reasonable shape, or you have just finished a build
- Your executive director, staff or funders are asking what you are doing about AI
- Staff already use AI tools, and nobody has decided the rules
- A pilot went nowhere, and you want to know why before trying again

It is probably not the right step if

- Your members or donors live in several systems. Start with the Digital Readiness Blueprint
- You want a chatbot live next week
- You have already chosen a use and need it built. Talk to us about that directly

In those cases we will say so, and point you to what fits better.

03 · HOW THE WEEKS RUN

Your team checks every finding before it becomes a plan

The same method as the Digital Readiness Blueprint, pointed at a different question.

WEEKS 1 TO 3

Inventory and goals

A kickoff to agree what leadership wants from AI and what is off limits. A short survey asks every team what they already use. Then a 45-minute interview with each team.

- Every AI tool in use, including features that arrived inside software you already pay for
- Ideas from membership, fundraising, programs, events and education, finance and communications

WEEKS 3 TO 4

Scoring, and your corrections

Each idea is scored by value, by risk and by how ready its data is. We score the data-readiness grid with the owner of each area. Then we read it all back.

- Your team corrects what we got wrong
- Your leadership ranks the portfolio, and the plan follows that ranking, not ours

WEEKS 4 TO 6

Policy, roadmap and the decision

The one-page AI policy, the 12-month roadmap and a decision briefing in your voice. The first use is specified down to its success number and stop rule.

- We walk you through all of it, and check in two weeks later
- Ready to start with us, with another partner, or on your own

What we need from you

45 minutes
with each team, in one conversation

A five-minute survey
for staff, about the tools they already use

An hour
with the owner of each area, to score the data together

An honest answer has to be safe

Staff almost always use tools nobody approved, and most of it is reasonable. We say up front that the aim is to approve what is working and agree a few lines, not to take tools away. People answer honestly when it is safe to, and the policy then governs what is really happening.

WHY YOUR TEAM RANKS THE PORTFOLIO

You know which work matters most and which mistakes your members or donors would not forgive. Every score in the portfolio has your team's hand on it, which is what makes the plan hold up when people ask why this use and not another.

How risk is set: by what happens when it is wrong

Low

A draft a person edits

Summaries, first drafts, inbox sorting. A person reads it before anything leaves.

Medium

An answer someone reads directly

A member or donor sees it without a person in between. Needs a test set and handover rules.

Keep a human

A decision about a person

Eligibility, scholarships, lapses, credit. AI can prepare the file. A named person decides.

04 · WHAT YOU RECEIVE

Seven things, every one of them yours to keep

Written for the people who will use them: your leadership, your staff, and whoever builds the first use.

DELIVERABLE	WHAT IT IS	YOU USE IT FOR
AI in use today	Every tool staff use, including personal accounts and features that arrived quietly, and what goes into each. Found in a way that makes an honest answer safe	Knowing what the policy has to govern
Use-case portfolio	15 to 30 uses from across the organization, each scored by value, risk and data readiness, and ranked by your leadership	Choosing what to start, and what to leave alone
AI data-readiness grid	Six rows, scored with the owner of each area, for the records each top use would read: what is ready, and what to fix first	Knowing which uses must wait, and for what
One-page AI policy	A draft ready for your leadership to adopt: four tiers of information, three review levels, disclosure, measures and a switch-off	Rules anyone can read in five minutes
12-month AI roadmap	Now, Next and Later. Each use with an owner, a number, and the data work it waits on	Funding and staffing in stages
Decision briefing	Eight slides in your voice, with speaker notes, ready for your board: what staff use today, the portfolio, the policy and one clear decision	The meeting where it is decided
The first use, specified	What it does, what it can see, what it cannot do, who reviews it, the number it should move and the rule for stopping it	Building it, with us or anyone else

Vendor-neutral, in practice. We recommend the kind of tool each use needs and the terms it should run under, not a vendor we resell. Some uses fit the AI already inside software you pay for; some fit another model better. You get the reasons, and ten questions to ask any vendor.

Do we need to be on Salesforce?

No. The inventory, the portfolio, the policy and the roadmap work whatever systems you run. Where a use would read member or donor data, the grid shows what that record needs, wherever it lives.

Had a Digital Readiness Blueprint?

Every one includes a short AI chapter: which uses could start early, and which to leave alone. The AI Blueprint goes the whole way.

05 · WHAT IT LOOKS LIKE

Excerpts from an AI Blueprint

Lakeview Association · example

11 tools	4	24 uses	5
already in use	on personal accounts	scored across eight areas	to start now

From the use-case portfolio

USE	AREA	RISK	DATA	SCORE
Renewal risk	Membership	Low	Ready	9.1
Donor next step	Fundraising	Low	Not yet	8.8
Monthly report	Leadership	Low	Ready	8.7
Member answers	Events and education	Medium	Not yet	8.4

From AI in use today

A renewal list went into a chat assistant on a personal account, to tidy the names. Nobody did anything wrong by the rules, because there were none. The policy now says which tools may see member data.

Fix first

From the data-readiness grid

Renewal risk: 10 of 12. One record, consistent fields and known owners are in place. Ready for a small, measured pilot.

Donor next step scores high and still waits: giving history sits in two places.

From the roadmap

NOW, MONTHS 1 TO 3
Renewal risk, monthly report, inbox triage, CE certificates, duplicates

NEXT, MONTHS 4 TO 6
Six uses, including a staff assistant and grant reports with figures checked

LATER, MONTHS 7 TO 12
Six uses, including member answers after a test set

WAITING ON DATA
Seven uses, including donor next step, after one record

From the first use, specified: UC-01, renewal risk list

WHAT IT DOES	Flags members drifting away each week, with the signals behind each, and drafts a check-in for staff
CANNOT	Send anything, or change a membership status
REVIEW	The membership manager reads every draft
SUCCESS NUMBER	Renewals among flagged members, against the same months last year
STOP RULE	No movement after one full renewal cycle: switch it off, and say why

06 · HOW ORGANIZATIONS USE IT

To set the rules

The policy is adopted, so staff know which tools they may use, what those tools may see and who reviews the result.

To start small

One use is approved with its number and stop rule. The next two come back for a decision with results in hand.

With whoever builds

The specification and the grid go to any partner, so the first build is scoped from your needs, not their product.

Questions people ask us

How is it different from the Digital Readiness Blueprint?

Same method, different question. That one decides what systems you need. This one decides where AI should work for you.

Which should come first?

If your members or donors live in more than one system, the Digital Readiness Blueprint. AI on a split record repeats the split, faster. If the record is in reasonable shape, start here.

How much of our team's time does it take?

A kickoff call, a five-minute survey, one 45-minute conversation per team, an hour with each area owner, the readout and the walkthrough. For most people, a few hours across the four to six weeks.

Do we have to use Kloudserv afterwards?

No. You own everything we deliver. Start the first use with us, with another partner, or on your own. We check in two weeks later to answer questions.

What does it cost?

A fixed fee, quoted before you commit.

What if we disagree with a score?

Good. The readout exists for that. Corrections are folded in, and your leadership ranks the portfolio before any roadmap is written.

Two Blueprints, one method

	DIGITAL READINESS BLUEPRINT	THE AI BLUEPRINT
The question	What systems do we need, and in what order?	Where should AI work for us, and what has to be true first?
Best for	Disconnected systems, a system to replace, or a move to Salesforce	Systems in reasonable shape, or a build just finished, and AI questions arriving
Both	Interviews with each team, your team corrects the findings, a decision deck or briefing in your voice, a fixed fee quoted before you commit, four to six weeks, vendor-neutral, yours to keep	

WONDERING WHERE AI FITS?

A 30-minute conversation is the first step

What your staff already use, where AI could help, and which Blueprint should come first. We will tell you honestly if you need neither.

Talk it through with us